Volume 32 Number 1
FEBRUARY/MARCH 2018

The Journal of
Information Integration and Management

database
TRENDS AND APPLICATIONS

The New Database Technology Landscape

4

WWW.DBTA.COM

The Journal of Information Integration and Management

database
TRENDS AND APPLICATIONS

Volume 32 | Number 1
FEBRUARY/MARCH 2018

CONTENTS

PUBLISHED BY Unisphere Media—a Division of Information Today, Inc.
EDITORIAL & SALES OFFICE 121 Chanlon Road, New Providence, NJ 07974
CORPORATE HEADQUARTERS 143 Old Marlton Pike, Medford, NJ 08055

DATABASE TRENDS AND APPLICATIONS (ISSN: 1547-9897; USPS: 16230) is published bimonthly (Feb./Mar., Apr./May, Jun./Jul., Aug./Sep., Oct./Nov., and Dec./Jan.) by Unisphere Media, a division of Information Today, Inc., 143 Old Marlton Pike, Medford, NJ 08055 USA; Phone (609) 654-6266; Fax (609) 654-4309; Internet: infotoday.com. Registered in U.S. Patent & Trademark Office. Periodicals postage paid at Vincentown, NJ, and additional mailing offices.

FEATURES

DEPARTMENTS

COLUMNS

MEDIA PARTNER OF THE FOLLOWING USER GROUPS

Citizen Developers Are Encouraged and Are Delivering Value

By Joe McKendrick

"Citizen developers"—those business users outside of the IT department who design or build their own applications—are more than a small band of rebels. A majority of organizations now rely on non-IT developers for at least some of their mission-critical applications. Few companies discourage such activity, and many are benefiting from the faster pace of software releases with which citizen developers play a part.

That's the gist of a recent survey of 324 executives conducted by Unisphere Research, a division of Information Today, Inc., publisher of *DBTA*. Respondents came from organizations of all sizes and across various industries. The report, titled "The Rise of the Empowered Citizen Developer: 2017 Low-Code Adoption Survey," was based on a survey of readers of *KMWorld* and *CRM* magazines and sponsored by Kintone. The majority of respondents, 82%, identified themselves as coming from outside the IT department.

The need for empowered business users is urgent. In today's tech-driven economy, IT has become the lifeblood of enterprises. At the same time, it requires many people across organizations to be able to design and build applications that help move information to the places where it is needed. In essence, IT has become much bigger than the IT professional or developer—IT has become the business.

Citizen developers are already a key part of most enterprises. At least 76% of respondents indicated that some portion of their applications were developed outside of their traditional IT department or IT service (see Figure 1). Executives and their staffs have some programming skills, but more than one in four know nothing about programming. Still, a majority have downloaded applications on their own, and close to half have worked directly on corporate websites or mobile apps.

Non-IT developers come from a range of backgrounds, but are, for the most part, power users and developers embedded within line-of-business departments building the applications. Outside consultants and line-of-business employees partake in much of this activity as well. For the most part, this off-the-grid IT work takes place on company hours, suggesting that citizen developers are accepted within the workflow. Close to half, 45%, report that all outside IT work is conducted during regular company hours.

Figure 1: Percent of Applications in Use Developed Outside of IT Department/IT Service

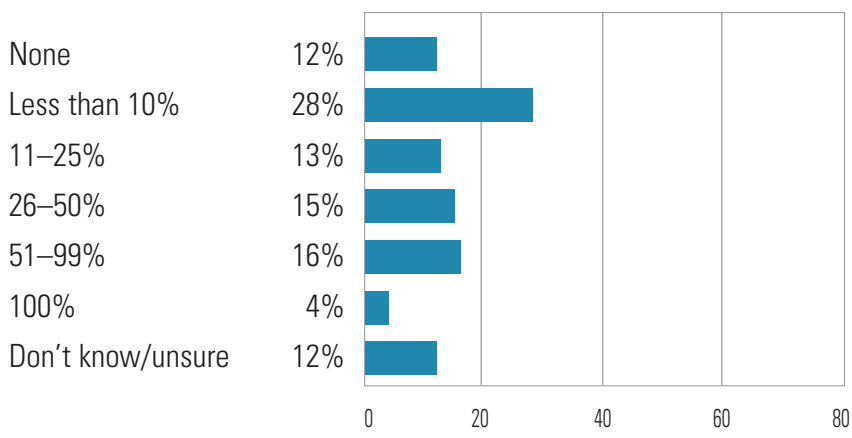

Figure 2: Level of Dissatisfaction With IT Department/IT Service (Reporting 1 or 2 on a scale of 1 to 5)

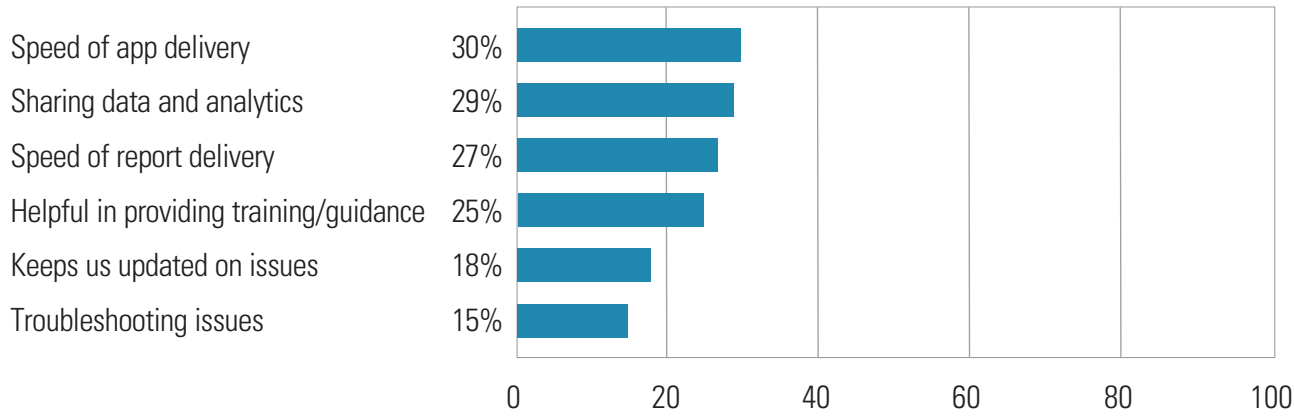

It's not that organizations are lacking in IT departments or trained IT talent. Most have some type of IT department staffed with at least one dedicated employee. A majority of executives, 74%, say their IT department is actively engaged to some degree, overseeing most technology-related activities.

What is driving the rise of the citizen developer? Speed of application delivery and the sharing of data and analytics are two areas in which IT support is seen as weak. Close to one-third, in fact, gave low marks to their IT departments in terms of timely delivery of software (see Figure 2). Citizen developers get applications out the door faster than large IT departments. They turn around their required applications in a matter of weeks, or a couple of months. Only 17% report turnaround times exceeding 3 months.

Citizen developers do what they do because they feel IT departments—which are usually weighed down with more critical responsibilities to keep enterprise applications up and running and secure—are too slow to respond to their individual requests. Recognizing that their IT departments are stretched to the limits, and need to focus on the big things—such as working with the overall business on technology strategy, as well as keeping backend systems humming—most organizations are on board with non-IT software creation. Only 16% attempt to clamp down on citizen development activity—more than one in four have no policy of any kind in place, while 42% say non-IT app development is allowed, or in some cases, actively encouraged.

Almost all executives acknowledge more needs to be done to provide training and support to citizen developers. However, only one-third of organizations are considered to be highly proactive in supporting their citizen developers with training and platforms.

Three types of resources have facilitated the rise of citizen developers: the rise of cloud, mobile, and open source solutions. The plethora of open source projects and offerings now available offer a wealth of possibilities for the citizen developer. A majority, 54%, turn to open source software as their first choice in building and supporting their self-created applications.

Of course, having citizen developers in the enterprise is not without its difficulties. The challenges to citizen development include issues around data security and trouble learning proper programming techniques and handling of data. ■

To download this research report, go to www.dbta.com/DBTA-Downloads/ResearchReports/THE-RISE-OF-THE-EMPOWERED-CITIZEN-DEVELOPER-7575.aspx


8:20
www.

The New Database Technology Landscape

From Relational to Blockchain

By Joe McKendrick

AN ASTOUNDING ARRAY of new technologies and approaches have emerged on the database scene over the past few years that promise to turn the next 12 months into a time of unprecedented transformation for the database landscape. There are new developments, along with reinforcement of tried-and-true technologies, some of which may help make the jobs of data managers just a bit easier. ▸▸

"Gone are the days of a terabyte of data sitting in a relational database accessed by a few analysts using BI tools," said Chris Doolittle, principal consultant of Teleran. "Big data, IoT, specialized database platforms, AI and machine learning, and the cloud are driving a generational transformation in data management."

There is still plenty of hard, brain-twisting, arm-twisting work ahead to get this next generation of technologies into and aligned with organizations. "We are on the cusp of an unprecedented intelligence revolution, and a lot of the enabling technologies—cloud, machine learning, artificial intelligence, real time databases, next-generation memory technologies—are already available," said Leena Joshi, VP of product marketing at Redis Labs. "What is needed is for enterprises to develop stacks that can tie all the piece parts together without generating layers of additional complexity." This, more than anything, describes the job of data managers in the year 2018.

Here are key developments that need to top data managers' to-do lists in terms of technology focus this year:

ANALYTICS WITH A PURPOSE

For a number of years, the goal of many enterprises—egged on by vendors and analyst groups—was to find ways to disperse analytics across the enterprise, a kind of "data democracy." Now, it may be time to shift gears on this vision, employing analytics not to empower single individuals, but to build a collaborative culture. "The trend toward self-service analytics is not panning out," said Jon Pilkington, chief product officer at Datawatch. "Putting analytics power in the hands of the business user was supposed to create agile companies and deliver analytical, data-driven decisions. Instead, companies are in worse shape than ever before. IT has lost control over data usage, and analysts are working in silos, duplicating work efforts and experiencing a severe lack of trust in their data and analytics outcomes."

Pilkington urges data managers to move away from the self-service goal and work toward more collaborative "team-based, enterprise data preparation and analytics." Such collaboration "will create a data-driven culture by bringing analysts together for the common purpose of getting answers—answers that are founded in the cross-business insights necessary to profoundly impact operational processes and the bottom line. Teams will be able to create, find, access, validate, and share governed, trustworthy datasets and models for true enterprise collaboration and faster, more strategic decision making."

ARTIFICIAL INTELLIGENCE TO IMPROVE ARTIFICIAL INTELLIGENCE

Artificial intelligence may go a long way in helping businesses understand and predict their futures, but AI is only as good as the data feeding it. Ironically, AI will help organizations achieve better AI results. "Many companies are faced with challenges around whether their data is current, complete, and consistent," said Doug Rybacki, VP of product management at Conga. "When you apply intelligent tools against data that is lacking in these components of data quality, the result is disappointing and potentially misleading." Ideally, said Rybacki, intelligent tools must be used for data hygiene so that the larger benefits of machine learning and artificial intelligence applications can be realized.

AI and associated machine learning will play a role in preparing and cleansing the data needed to make AI and machine learning work. "We'll see more adoption of machine learning to assist with all aspects of data fabric and data management," said Tendu Yogurtcu, CTO of Syncsort. "This will span data integration to data quality, as insights are only as good as the data that drives them, and only useful if the data is error-free and ready for advanced analytics." As proof of this need, Yogurtcu noted that "data scientists are still spending more than 80% of their time on data preparation. Bringing machine learning into data management processes will help ensure automation of these data preparation steps. Using machine learning to drive business rules, in the data cleansing and modeling processes, will free up data stewards and data scientists to focus on deriving actionable insights from the data."

There's even a name for this application of AI to improve AI. Jerry Melnick, president and CEO of SIOS, sees "AIOps"—which is the use of machine learning applications in IT—as a growing proposition. Most IT teams "are broken up into silos, and each silo has its own set of analytics and diagnostic tools it uses to trace performance issues," said Melnick. "AIOps eliminates this issue, using machine learning to track the relationship between every element of IT, and understand how these elements interact with each other. This elevates problem-solving about the siloed webs of IT, and gives teams actionable, data-driven solutions to IT's biggest headaches."

CompuCom CIO and CDO Justin Mennen agreed that AI and associated technologies will significantly shape data management practices, especially "when we consider the overwhelming need for organizations to leverage predictive and prescriptive analytics to compete." Progress on this front, however, "requires a new construct beyond the data mining BI model of the last 2 decades," Mennen said. He urges looking at data in new ways, including "the use of data transformation and graph models with analytics to view the intersections, correlations, and isolations among the noise," combined with "the communication and storytelling of potential wisdom."

Ultimately, of course, AI needs to deliver to the business. Doolittle sees more intelligent data management solutions that "combine machine learning with rule-based systems to watch and learn from changing data usage patterns and user behaviors. They automatically create data management rules that can automatically direct changes or actions to better serve changing business demands. Examples include identifying resource-consuming user behaviors that indicate a need for more shifting data workloads to a more appropriate or cost-effective data platforms, or increasing use of more detailed data indicating a demand for direct access to source data to improve analytical outcomes and lower data handling costs."

VIRTUAL ASSISTANTS

Another technology development that is fueled by data is the rise of virtual assistants. The cutting-edge web companies are employing this type of solution, and it is coming to mainstream enterprises as well. Google, for example, is thriving with its Google Now virtual assistant, "which is only getting smarter because of its ability to use available data from web interactions to provide a personal experience for users," observed Luc Burgelman, CEO of NGDATA.

"Companies—especially those in customer-facing industries such as banks, media, and retail—need to engage with and support customers through conversational interfaces, so we'll see them add more artificial intelligence and cognitive services to their offerings to create interactive experiences," said Burgelman. Data and the effective management of it are at the heart of such capabilities, he added.

The most important part of creating virtual assistants will be the data and having the data drive actions and decisions, Burgelman noted. "This means considering all data—including real-time and behavioral data—and learning from all channels to create connected experiences customers expect. Powering these customer interactions through the understanding of all this detail will be critical for companies."

CLOUD ADVANCES

Cloud computing, which has been a major force in the IT and data management space for close to a decade, continues to reshape database technologies as well. Cloud is increasingly the home of "systems of insight" that support advanced data analytics and artificial intelligence capabilities, said Roman Stanek, CEO and founder at GoodData. In Stanek's view, a "unified technology platform in the cloud is the future of analytics and data in the cloud." Data in the cloud is growing rapidly, and there is no way to manage that other than through a system of insight, he added. The industry is facing a confluence of trends, he noted. These include data growing exponentially and old BI failing, while advances in BI such as machine learning and predictive analytics make it ripe to take off.

Not only will there be systems of insight in the cloud, but multiple clouds for multiple use cases as well. Lately, there's been movement to multi-cloud strategies, especially as more applications and innovations open up. "Most companies don't set out to adopt a multi-cloud strategy," said Jaspreet Singh, CEO and founder of Druva. "Rather, they choose to work with cloud vendors for specific use cases, and when we take a step back, we see a multi-cloud implementation. In that regard, multi-cloud is not a strategy, it's an outcome of these decisions."

BRING ON THE BLOCKCHAINS

Another technology that is seriously being explored by many enterprises is blockchain—an online global database that stores and manages smart contracts and transactions. Some observers see blockchain as the next great frontier for data management. "Right now, we only hear about blockchain with cryptocurrency and are starting to see emerging companies in the finance and healthcare space discussing the value," said Avani Desai, executive VP of Schellman and Company.

The value may be in blockchain's distributed nature—data is verified across multiple nodes, and thus protected from tampering. "Blockchain provides you 100% assurance that a transaction was valid," Desai observed. "It shows me what was done, by whom, when, and maybe even the why. This provides transparency and reconciliation, one of the most difficult aspects of a distributed system."

LATENCY BUSTERS

The move toward real-time computing and real-time enterprises is also shaping the database technology landscape this year. At the same time, many of the technologies with which data management teams are working may add more latency into transactions and computing jobs. "While moving to real-time is a trend, it competes directly with the move to microservices, distributed logs, and asynchronous stream processing," said John Hugg, founding engineer and manager of developer outreach at VoltDB. "All of these things can make our systems more resilient and dynamic, but they often compound latency problems. Things that used to be a single network round trip might become dozens of asynchronous network messages." ▶▶

The move to real time also poses challenges for data quality, Hugg continued. "Immediacy and correctness are often difficult to manage at the same time. Many systems do a poor job of processing late and out-of-order data, preferring to collect mini-batches and process them whole. When latency is paramount, you see a return to the simpler request/response model of traditional services." This requires technologies that can step up and compensate for these new sources of latency, such as in-memory caches, grids, and databases enabling real-time processing, Hugg explained.

Often, even cloud, with all its unlimited power, may not be the right fit for organizations seeking to get closer to real-time movement of information. For example, while the "movement of in-memory databases is fundamental to the applications, the truth is that deploying the infrastructure to support these uses outside of a data center or exclusively in a cloud is hard for real-world situations," said Jason Andersen, VP of business line management at Stratus Technologies. "There is a need for a better infrastructure to protect real-time data at these edge locations."

Nowhere is the need for real-time and reduced latency felt more strongly than in efforts to leverage the Internet of Things (IoT). Capturing data in real time, tied to IoT, can be effective only with systems capable of cost-effectively handling large data volumes with very low latencies, said Joshi of Redis Labs. "Being able to implement adaptive applications powered by machine learning in real time is a critical aspiration for most enterprises, but real-time databases that can power such applications with built-in capabilities are most likely to make these aspirations a reality." Joshi added that another critical force in making the data-driven enterprise a reality is the shift in hardware technology, which puts more cost-effective memory such as flash within reach of applications. "Datasets that can deliver the real-time performance of memory but with the cost-effectiveness of flash are likely to create a competitive edge for enterprise," she said.

METADATA AND DATA CATALOGS

Despite all the hype and excitement about data-driven, AI-savvy enterprises, there is a fundamental component of data management that managers are beginning to embrace: keeping track of data assets and making them discoverable to decision makers. With data streaming in from a wide variety of internal and external sources, there needs to be a way to intelligently track, archive, and identify what is available to decision makers and applications. Metadata repositories and data catalogs are the way this can be achieved. "People tend to focus on things like in-memory and other speed-and-feeds sorts of metrics when they think about real-time technologies," said Joe Pasqua, executive VP at MarkLogic. "But that assumes you've got all the relevant data in one place and you're just trying to serve it up quickly. That's the easy part. The real enabler is making the data available in the first place. This is made possible by a strong metadata solution to describe what and where the data is, and a multi-model approach that allows access to the varied shapes, sizes, and formats of the data across your organization, including graphs, documents, rows, geospatial, and so on."

Metadata is also key to the success of AI, as well. "When AI can be leveraged to automatically and accurately append metadata attributes to information, the whole game changes," said Greg Milliken, senior VP of marketing for M-Files. "AI can automatically evaluate the file contents for specific terms like a customer name, project, or case as well as the type or class of document—a contract, invoice, project plan, financial report—and then apply those metadata attributes to the file. This automatically initiates workflow processes and establishes access permissions, so only authorized people can access the information—such as the project team, the HR department, or those managing a specific case." The result, Milliken continued, "is a more intelligent and streamlined information environment that not only ensures consistency in how content is organized, but also that information is intelligently linked to other relevant data, content, and processes to deliver a 360-degree view of structured data and unstructured content residing in different business systems."

OPEN SOURCE PREVAILS

Open source technologies have emerged that support the emerging real-time data center, said Marc Concannon, CTO of Clavis Insight. These include Kafka for capturing and distributing incoming streaming data; NiFi for data routing; Ignite for faster in memory process of the incoming data; Hadoop 2.0 for data access and storage; and Kubernetes for managing how we scale a streaming infrastructure which is susceptible to bursts. "All of these technologies are relatively new to our stack," Concannon pointed out. "But these technologies at their core are all about working with more and more data and extracting the relevant insights from this data quicker and hence making it available to our customers quicker."

Hadoop "does enable a lot of tools which are focused on streaming and it also enables quicker access to the core insights on large datasets which is not really possible or would mean a long wait on more traditional technologies, said Concannon. "This, to me, is all about making more data available at decision time." ■

How to Maximize Your Company's Location Data

By Javier de la Torre

Download any app onto a smartphone these days and you'll be prompted by the same question: Do you want to turn location services on?

For consumers, letting Google Maps or Waze know exactly where you are is a must—it's the essence of why those apps exist. But have you ever wondered why other types of businesses, from grocery stores to social media platforms, want to know where their patrons are at all times?

It turns out this across-the-board drive from enterprises is not unique. According to a study by Dresner Advisory Services, more than 60% of people realize that location intelligence (LI) is a critical component of business.

So what makes LI unique? The truth is that almost every business has a treasure trove of location data—for example like Starbucks knowing exactly where everyone with their app is right now. However, the business intelligence (BI) tools being used to process this data can only provide temporal insight or basic points on a map. LI goes a step further by answering more forward-thinking questions only possible with an LI platform. Answering, for example,"Where should Starbucks put its stores in the future based on user commuting patterns?" Or, "Where would a billboard reach the largest number of prospective buyers early in the morning?" Instead of focusing on the now, LI seeks to answer challenges a business may face moving forward.

Visualizations are the foundation for interpreting LI—the baseline for how you will better understand your data.

But performing this type of iterative spatial analysis involves so much more than just collecting the data and hoping it points an enterprise in the right direction. To get the most out of location-based services—now ubiquitous thanks to the explosion of connected devices, the Internet of Things, mobile phones, connected cars, and social media—companies should follow a few best practices so they don't get lost on their way to LI-driven problem solving. ▶▶

APPLICATIONS ▸▸

STEP 1: Enrich the Data

Having quality data is the first step to gaining LI insights. Companies must ensure their data is accurate and reliable, cleaning and filtering it before integrating it with other external databases. Any LI services providers should be well-versed in how to get only the best data imported or connected to a new database. Then, it's time to enrich it—linking it with additional sources of location data, such as financial, ecological, and demographic measurements. To determine which sources of data to add, company executives should ask themselves what they want to achieve, where their data is located, and what's the best way to aggregate their data.

STEP 2: Create a Visualization

Visualizations are the foundation for interpreting LI—the baseline for how you will better understand your data. There are multiple methods to selecting how to best chart your data. Some things to consider are how much you'll need to interact with your data to model scenarios and findings correctly, if the visualization is easy to understand (particularly important if the visualization will be interpreted by a team), and if it appropriately represents the data and findings?

Think about how the National Weather Service recently had to redesign Hurricane Harvey rain data on-the-fly, so it could be more easily interpreted by the public. This kind of flexibility and ease of use, where a tool scales with immediate needs, is essential.

STEP 3: Analyze and Iterate

This step is what separates BI from LI. Your initial analysis, or blend of analysis methods, should be based on your company's desired outcomes, which can come in a broad array of functionalities. By combining traditional analytics with emerging geospatial techniques, companies can create more actionable business outcomes and value from their data.

Spatial data analysis marries scientific modeling and machine learning to perform actions like identifying clusters and outliers and predicting market volatility or future consumer patterns. Database analysis works best for filtering, numeric aggregations, joins, and other methods traditionally also found in BI tools. Geospatial analysis is practical for any spatial functionality, such as measuring distance, proximity, or other functions often found in geographic information system (GIS) platforms.

Once you select one or a few methods, your LI dashboard should enable your team to iterate on this data. After collaborating, an organization's teams can then move onto the final step.

STEP 4: Take Action

Now that you have your findings, communicating them with other stakeholders and determining a course of action are imperative to turning your LI data into concrete changes to your future business practices. Companies may even opt to develop new applications based on what they've learned, creating new product opportunities and streams of revenue. Or companies can identify how to increase their profit margins based on their LI insights.

The Result

Collecting LI data is just the tip of the iceberg when it comes to using that data to drive insights. Using location-based information, businesses could shift store hours to reflect when customers are actually around their business. They could gauge real estate market pricing based on projected foot traffic. They could perform geomarketing to ensure ads are placed in the right region based on demographics, and businesses can make all of these choices knowing they aren't just going on a hunch. Each of these scenarios can lead to better outcomes that would have never been possible without LI best practices. ■

Javier de la Torre is CEO of CARTO.

United States Postal Service
STATEMENT OF OWNERSHIP, MANAGEMENT AND CIRCULATION (Requester Publications Only)

1. Publication Title: Database Trends and Applications. 2. Publication Number. 16-230. 3. Filing Date: 10/1/17. 4. Issue Frequency: Bi-Monthly; Dec/Jan., Feb/Mar., Apr/May, June/Jul., Aug/Sept., Oct/Nov. 5. Number of issues Published Annually: 6. 6. Annual Subscription Price: 0. 7. Complete Mailing Address of Known Office of Publication: Information Today, Inc., 143 Old Marlton Pike, Medford, Burlington County, NJ 08055. 8. Complete Mailing Address of Headquarters or General Business Office of Publisher: Unisphere Media a Division of Information Today, Inc., 143 Old Marlton Pike, Medford, NJ 08055. 9. Full Names and Complete Mailing Addresses of Publisher, Editor, and Managing Editor: Publisher: Thomas Hogan, Jr. Group Publisher, 143 Old Marlton Pike, Medford, NJ 08055-8750. Editor: None. Managing Editor: Joyce Wells, 121 Chanlon Road, New Providence, NJ 07974. 10. Owner: Information Today, Inc., 143 Old Marlton Pike, Medford, NJ 08055, Thomas H. Hogan, 143 Old Marlton Pike, Medford, NJ 08055, Roger R. Bilboul, 22 Earls Terrace, London W8 6LP, England. 11. Known Bondholders, Mortgagees, and Other Security Holders Owning or Holding 1 percent or More of Total Amount of Bonds, Mortgages or Other Securities If none, check box: None. 12. Has Not Changed. 13. Publication Title: Database Trends and Applications. 14. Issue Date for Circulation Data Below: Oct/Nov. 2017. 15. Extent and Nature of Circulation: a. Total Number of Copies (Net press run): Average No. Copies Each Issue During Preceding 12 Months, 13,099; No.Copies of Single Issue Published Nearest to Filing Date; 9650. b. Legitimate Paid and/or Requested Distribution. (1) Outside-County Paid/Requested Mail Subscriptions Stated on Form 3541 (Include direct written request from recipient, telemarketing and Internet requests from recipient, paid subscriptions including nominal rate subscriptions, employer requests, advertiser's proof copies, and exchange copies): Average No. Copies Each Issue During Preceding 12 Months, 7652; No.Copies of Single Issue Published Nearest to Filing Date, 4843. (2) In-County Paid/Requested Mail SubscriptionsStated on Form 3541 (Include direct written request from recipient, telemarketing and Internet requests from recipient, paid subscriptions including nominal rate subscriptions, employer requests, advertiser's proof copies, and exchange copies): Average No. Copies Each Issue During Preceding 12 Months, 0; No.Copies of Single Issue Published Nearest to Filing Date, 0. (3) Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Paid or Requested Distribution Outside USPS: Average No. Copies Each Issue During Preceding 12 Months, 0; No.Copies of Single Issue Published Nearest to Filing Date, 0. (4) Requested Copies Distributed by Other Mail Classes Through the USPS (e.g., First Class Mail): Average No. Copies Each Issue During Preceding 12 Months, 0 ; No. Copies of Single Issue Published Nearest to Filing Date, 0. c. Total Paid and/or Requested Circulation (Sum of 15b (1), (2), (3), and (4): Average No. Copies Each Issue During Preceding 12 Months, 7652; No.Copies of Single Issue Published Nearest to Filing Date, 4843. d. Non-requested Distribution (By Mail and Outside the Mail) (1) Outside County Non-requested Copies Stated on PS Form 3541. (Include Sample copies, Requests Over 3 years old, Requests induced by a Premium, Bulk Sales and Requests including Association Requests, Names obtained from Business Directories, Lists, and other sources): Average No. Copies Each Issue During Preceding 12 Months, 5137; No.Copies of Single Issue Published Nearest to Filing Date, 4617. (2) In-County Non-requested Copies Stated on PS Form 3541 (include Sample copies, Requests Over 3 years old, Requests induced by a Premium, Bulk Sales and Requests including Association Requests, Names obtained from Business Directories, Lists, and other sources): Average No. Copies Each Issue During Preceding 12 Months, 0; No. Copies of Single Issue Published Nearest to Filing Date, 0; (3) Non-requested Copies Distributed Through the USPS by Other Classes of Mail (e.g., First-Class Mail, Non-requestor Copies mailed in excess of 10% Limit mailed at Standard Mail or Package Services Rates): Average No. CopiesEach issue During Preceding 12 Months, 0; No. Copies of Single issue Published Nearest Filing Date, 0. (4) Non-requested Copies Distributed Outside the Mail (Include Pickup Stands, Trade Shows, Showrooms and Other Sources): Average No. Copies Each Issue During Preceding 12 Months, 96; No.Copies of Single Issue Published Nearest to Filing Date, 25. e. Total Non-requested Distribution (Sum of 15d (1),(2), and (3)): Average No. Copies Each Issue During Preceding 12 Months, 5233; No.Copies of Single Issue Published Nearest to Filing Date; 4642. f. Total Distribution (Sum of 15c and e): Average No. Copies Each Issue During Preceding 12 Months, 12,885; No.Copies of Single Issue Published Nearest to File Date; 9485. g. Copies not Distributed: Average No. Copies Each issue During Preceding 12 Months, 214; No. Copies of Single Issue Published Nearest to Filing Date, 165. h. Total (Sum of 15f and g): Average No. Copies Each Issue During Preceding 12 Months, 13,099; No.Copies of Single Issue Published Nearest to Filing Date, 9650. i. Percent Paid and/or Requested Circulation (15c divided by f times 100): Average No. Copies Each Issue During Preceding 12 Months, 59.39%; No.Copies of Single Issue Published Nearest to Filing Date, 51.06%. 16. Publication of Statement of Ownership for a Requestor Publication is required and will be printed in the December 2017/January 2018 issue of this publication. 17. Signature and Title of Editor, Publisher, Business Manager or Owner: John C. Yersak, Vice President & CAO. Date: 10/1/17. I certify that all information furnished on this form is true and complete. I understand that anyone who furnishes false or misleading information on this form or who omits material or information requested on the form may be subject to criminal sanctions (including fines and imprisonment) and/or civil sanctions (including civil penalties).

database
TRENDS AND APPLICATIONS

database
TRENDS AND APPLICATIONS
ONE COMPLETE MARKETING PROGRAM

Managing the HYBRID FUTURE

Best Practices Series

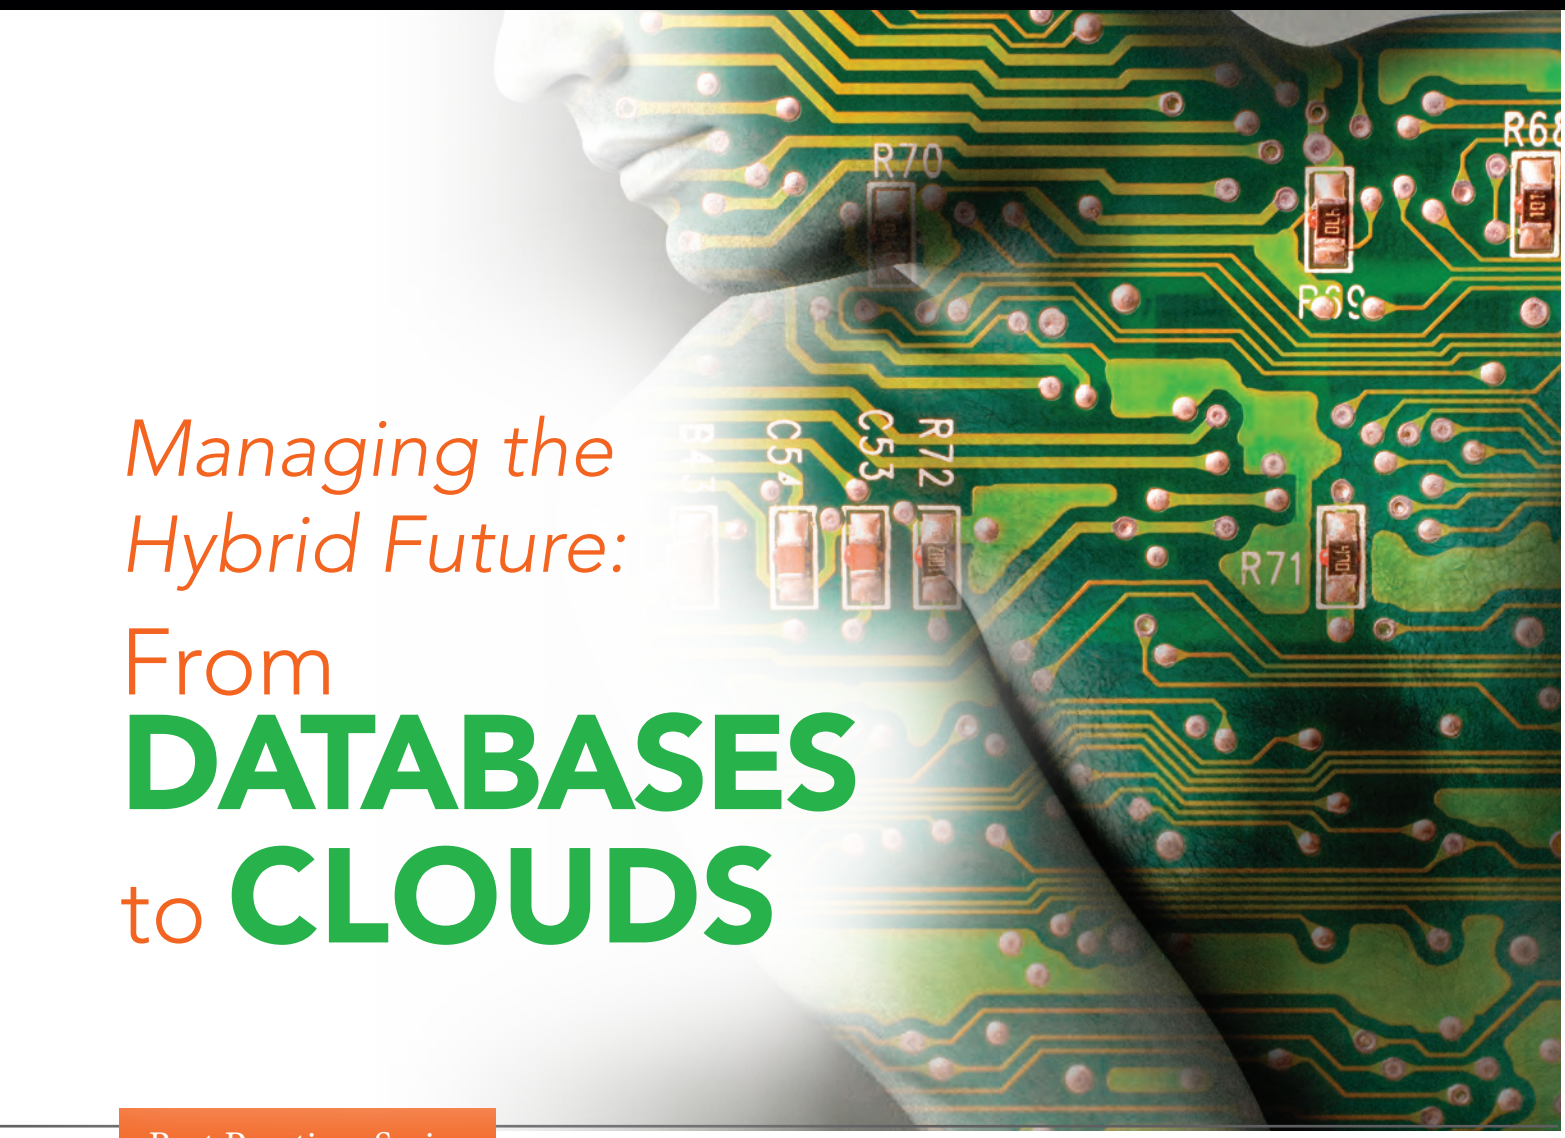

Managing the Hybrid Future: From DATABASES to CLOUDS

Best Practices Series

THE FUTURE IS IN THE CLOUD; the future is on-premises. There's been a steady drumbeat of opinion from vendors, analysts, conference presenters, and the trade press about an impending cloud future, in which any and all technology assets and systems will eventually end up in the cloud. While there's no question that cloud will play a prominent role in technology decisions and implementations, it is likely that enterprises will continue to need on-premises technology as well.

Cloud computing helps organizations move light years ahead into the digital realm, but many enterprises still need to commit much of their IT resources to existing on-premises "systems of record," a recent study conducted by Unisphere Research, a division of Information Today, Inc., shows.

The balancing act between managing cloud and on-premises environments is often a hit-or-miss proposition. The survey of 474 executives and professionals found that most organizational budgets still go to supporting existing enterprise systems and will continue to do so for the long-term future. On average, 45% of IT budgets go to supporting legacy "systems of record," while 31% go to newer "systems of innovation, according to the report ("Behind Every Cloud Transformation, There's Well-Functioning IT: 2017 Information Technology Performance Excellence Survey," January 2018).

Moreover, a survey of more than 2,000 IT professionals by Intel Security found that hybrid cloud adoption grew threefold over a 1-year period, increasing from 19% to 57% of organizations surveyed. At the same time, within the coming year, 80% of all IT budgets will be focused on cloud solutions.

At close to half of database sites, cloud will be shaping development and deployment over the next 3 years, the Unisphere survey found. Close to one in four, 24%, of executives indicated that cloud will play a "robust and strategic role" in the development and implementation of their databases, and another 24% stated that cloud will play a limited role. Another 29% expect cloud to be a major part of their ERP rollouts, as do 28% for core enterprise application suites in general.

IT and database managers should prepare to be working within both on-premises and cloud implementations for the foreseeable future. In the process, they will be able to take

advantage of the best of both worlds, leveraging solutions for their business where they fit best.

The future is extremely hybrid, and here are some guidelines to get the best out of it:

DON'T JUMP TO CLOUD JUST BECAUSE IT'S CLOUD

Cloud may appear to be a compelling value proposition, and, yes, everyone is talking about it. IT and data managers may be forgiven for thinking that the whole world is going to the cloud. But, as the Unisphere Research data showed, that is not the case, at least not yet. Most enterprises still have on-premises systems and assets that remain valuable components of their organizations. Cloud has many compelling benefits, including relatively low monthly subscription rates, as well as opportunities for greater flexibility as instances can be spun up in a matter of minutes. However, these advantages should be carefully weighed against the value that current systems are delivering and the potential for disruption in moving to a cloud provider. There may be a need for training, and entire job roles may need to shift as well.

BE NEUTRAL, AND SERVE AS A BROKER OF ESSENTIAL BUSINESS AND TECHNOLOGY SERVICES

Perhaps the best stance IT and data managers can take is the role of a technology "Switzerland," remaining neutral and objective when it comes to technology decisions. That's because the best solutions may come out of the enterprise's data center, or they may be available through the cloud. Business leaders and end users may not be aware of the pros and cons of each, but technology managers can make it their role to provide the right direction, no matter where the appropriate servers and applications are based.

CONSIDER DATA SECURITY, AND WHERE IT IS MOST SECURE

Many consider the cloud to be a risky place to store their data, but the risk is not necessarily for the reasons they think. In many ways, cloud environments may be more secure from hackers than on-premises data centers, as cloud providers make it their business to keep up with the latest security tools, standards, and best practices. At the same time, with data in the hands of a cloud provider, it may be difficult to get data returned in a timely manner in the event a contract is terminated, or if the cloud provider goes out of business. Vendor lock-in also was top-of-mind for many respondents in the Unisphere Research survey (40%), an issue that also ties in with loss of control, as it could be a messy process to move systems in or out of cloud environments that have proprietary hooks.

The best policy in this regard may be to maintain copies of the data in two places—one on-premises, and one in the cloud.

Many assume that private cloud offers greater data security than public cloud services. However, many of the threats to data security come from within. A privileged user with backdoor access or a disgruntled employee may pose just as much risk to data security as the hackers from the other side of the world. The production data center at the core of the private cloud may be locked down and sealed tight, but sensitive data may be sent out to parts of the enterprise that aren't exercising as much security. Many of the cases of data loss have occurred as a result of contractors carelessly leaving in the back of a taxi or a vehicle thumb drives or laptops loaded with sensitive customer data from their clients that is later stolen.

DETERMINE LINES OF RESPONSIBILITY FOR CLOUD VERSUS ON-PREMISES SYSTEMS, APPLICATIONS, AND DATA

As found in the Unisphere survey, IT and data managers aren't always clear as to who assumes responsibility for the risks in moving applications or systems to the cloud. Any glitches or disruptions in service will likely be at least partially someone else's responsibility to identify, find, and fix. The survey also finds that issues with applications and systems aren't always immediately known to IT personnel, and adding a cloud vendor between them and customers may slow down the identification process even further.

IT and database managers should prepare to be working within both on-premises and cloud implementations for the foreseeable future. In the process, they will be able to take advantage of the best of both worlds, leveraging solutions for their business where they fit best. ■

—Joe McKendrick

The Impact of Artificial Intelligence in 2018: Seven Predictions

By Joyce Wells and Stephanie Simone

Artificial intelligence has been described both as a radical force for good and an enabler of evil whose risks are not yet fully understood. But is it just a lot of hype or a transformative technology poised for impact? Here, seven IT execs weigh in on the changes AI may bring in 2018.

1 The AI debate shifts from "Is it good or evil?" to "Is it ever going to be good enough?": If 2017 was the year where the warnings from Elon Musk and Stephen Hawking about the potential evil from AI clashed with predictions from Mark Zuckerberg and Bill Gates on its potential good, 2018 will be the year when the debate shifts to its practical utility. Much like other technologies that were lauded for their world-changing potential and then fizzled as the fog of the hype cleared, early adopters will find themselves disappointed by AI's obvious limits. The broader public—familiar with Alexa, Siri, and Google Home—will be similarly disillusioned as the experts acknowledge that there is only so much that AI will be able to do, and for really complex problems, a new paradigm will be needed.

— Michel Morvan of Cosmo Tech

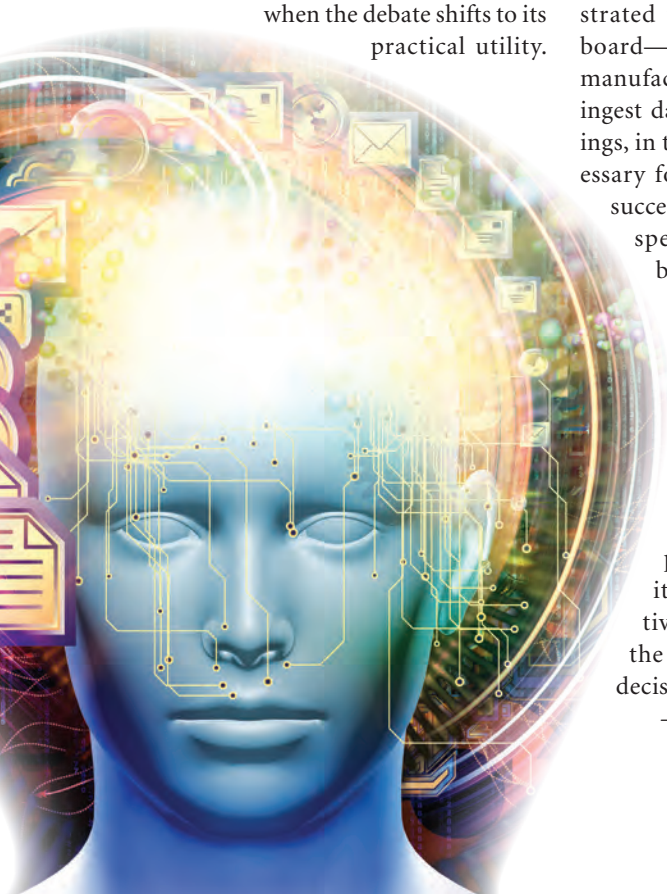

2 AI will guide us through the trees. Despite the hype, AI has demonstrated value in industries across the board—from agriculture to biotech to manufacturing. AI is just beginning to ingest data to power services and offerings, in turn providing information necessary for better decision making. AI's success will continue in the new year, specifically in a new area: troubleshooting. Expect to see an impact on troubleshooting for operators, data centers, etc., as AI helps individuals tackle the day-to-day issues, enabling them to focus on critical problems that AI itself can't help. In 2018, AI will guide and augment humans in solving hard problems as it further cements its value-add as a human cognitive partner, guiding us through the trees to make more impactful decisions.

— Ash Munshi, CEO, Pepperdata

3 Semantic Technology will become the AI Interpreter: As artificial intelligence becomes the new consumer-facing UI for many businesses, semantic technology will emerge as the necessary interpreter. Conversational AI will need precise understanding of the communication from humans and extract meaning from the communication. Artificial intelligence in combination with semantic technology is ideally suited to address this challenge.

— Jans Aasman, CEO of Franz

4 AI will be a creativity enabler: The role of the data analyst is changing thanks to artificial intelligence. AI is allowing marketers to focus once again on the creative art of marketing—the days of data wrangling are coming to an end. With studies indicating that up to 80% of an analyst's daily routine was relegated to data cleansing and preparation, 2018 will be the year where that 80/20 rule gets flipped upside down. The new AI-based approach to marketing technology will effectively create a win/win for both analysts and marketers alike.

— Leah Pope, CMO for Datorama

5 The button disappears and AI becomes the app interface: Developers need to figure out what data is really important to their business application, how to watch and learn from transactions, what business decisions would most benefit from this kind of proactive AI, and start experimenting. Embedded AI can predict what ▶▶

TRENDS ▸▸

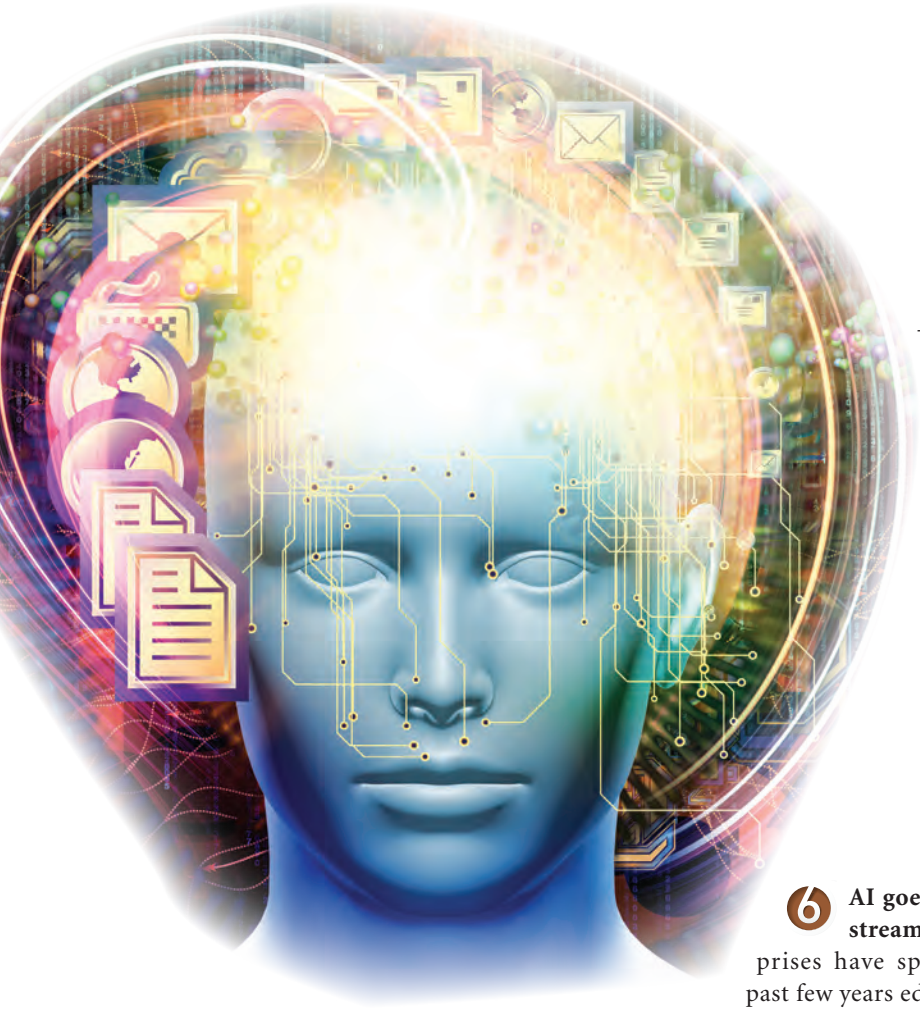

you need; deliver info and functionality via the right medium at the right place and time, including before you need it; and automate many tasks you do manually today. Some examples: An expense approvals app watches your pattern of approving expense reports, starts to auto-approve 99% of expense reports, and only brings to your attention the rare report that requires your attention; an analytics app understands the underlying data, questions asked so far by the business user, questions that other business users in the company might have asked of the same dataset, and each day provides a new insight that the analyst might not have thought of.

— Siddhartha Agarwal, VP, Product Management & Strategy, Oracle

6 AI goes mainstream: Enterprises have spent the past few years educating themselves on various artificial intelligence frameworks and tools. But as AI goes mainstream, it will move beyond just small scale experiments run by data scientists in an ad hoc manner to being automated and operationalized. The complexity of technologies used for data-driven machine and deep learning has meant that data scientists spend less time coding and building algorithms and more time configuring and administering databases and data management systems. And as enterprises move forward with operationalizing AI, they will look for products and tools to automate, manage and streamline the entire machine learning and deep learning lifecycle. Data scientists need to focus on the code and algorithms and not automating and operationalizing the process. In 2018, investments in AI lifecycle management will increase, and technologies that house the data and supervise the process will mature.

— Kinetica CTO and co-founder Nima Negahban

7 AI will not transform the enterprise in the near future: Previous predictions and claims about the direct impact of AI on enterprises have been overblown. There is excessive hype around how AI will lead us to new discoveries and medical breakthroughs. However, those expecting AI to be the ultimate truth conveyer are mistaken. It will be very hard to design a model that can determine unbiased truth, because human bias—whether explicitly or implicitly—will be coded into these data analytics systems and reinforce existing beliefs and prejudices. With that said, there are certain applications where systems can make better decisions in a shorter amount of time than humans, such as in the case of autonomous vehicles. In 2018 we will begin to see real use cases of the power of AI appear in our everyday lives—it just isn't ready to be the shining star for the enterprise quite yet. Only half of the Global 2000 offer fully digital products. So, despite all of the buzz around digital transformation, there's a lot of catch-up to be done before many of these companies can even consider looking at advanced developments such as AI.

— Christian Beedgen, CTO, Sumo Logic ■

MV SOLUTIONS

Revelation Software Enriches Riverside Cemetery With

DEEP RECORDS MANAGEMENT

By Stephanie Simone

Planned in the style of Frederick Law Olmsted's Central Park, Riverside Cemetery is a 100-acre facility that is more than a century old. Based in Saddlebrook, N.J., the company maintains public records for burials, private records on families, and billing. It also oversees all the work done to burial plots in addition to taking care of monuments and records relating to them, along with other tasks associated with running the business. The company has a strong focus on customer service and continuously modernizes operations to keep up with changing requirements.

"Most people think a cemetery is two people with shovels, people digging holes and putting caskets in the ground, but that's really not all a cemetery is—there is a lot happening behind the scenes," said Peter A. Blacksberg, president of Riverside Cemetery.

A cemetery is a combination of landscape work, record keeping, financial planning, accounting, and personnel management, as well as additional considerations such as maintenance vehicles, Blacksberg said.

At Riverside, all records are kept permanently and, until 25 years ago, all of it was stored on paper.

The 1980s brought dramatic changes to Riverside. Office procedures that had been stranded in the 1950s were replaced by early personal computers for letter writing and record keeping. Standards for each aspect of cemetery maintenance were established, and the field crews were professionally trained in each task.

This technological transition in 1980 was aided by WinWin Solutions, parent company of Revelation Software. Initially using Revelation Software's DOS-based Advanced Revelation, and then migrating to the Windows-based OpenInsight platform, Riverside has evolved technologically while remaining a traditional cemetery serving the community with warmth, compassion, and professionalism.

And today, Blacksberg believes Revelation's platform is at the center of an operations-based organization that no one ever sees.

*With Revelation Software's **OpenInsight platform, Riverside** evolved into a progressive powerhouse.*

Keeping It Straight With OpenInsight

Riverside Cemetery worked with Mike Ruane, president of WinWin Solutions, to overhaul document storage and maintenance procedures.

Riverside runs its business on OpenInsight, which helps with everything from keeping track of phone calls from funeral directors to printing and paying the bills and tracking employee hours.

"Nearly everything is structured through a process control system that we custom-designed with WinWin Solutions," Blacksberg said.

The entire workflow process is custom-built to meet the needs of the cemetery, Blacksberg explained.

"We have permanent relationships with families," Blacksberg said. "That's one of the many things that make our cemetery really different; we don't purge the customer list."

Using the database, the company records conversations with customers and can later pull up that data when customers call.

The computer system provides a very rapid and essential first view of what's going on, according to Blacksberg.

Transitioning Into the Future With Revelation Software

Changing from a paper environment to information automation is complex and not trivial, Blacksberg explained.

At first there was resistance to giving up doing work by hand. After 5–6 years of development and data entry, the cemetery was able to integrate more tools such as document merging, digital photography, and records-based digital mapping.

"Each year we process over 2,000 work orders on the cemetery grounds. With document merging, we were able to print a map of the precise location where the work is to be done. It saves time and limits frustration. We were able to move away from one-off word processor-typed letters to dynamic database merging," Blacksberg said.

Over time, the company has incorporated more tasks which have enabled them to use the staff in more effective ways.

Deciding on OpenInsight

Before choosing OpenInsight, Blacksberg created a prototype and educated himself on programming software to adjust and add features to the database. Blacksberg was already familiar with different methods, he had worked previously in Silicon Valley.

After being introduced to Revelation's software, he found the product to be malleable and adaptable to the cemetery's purposes.

"Even though we've used it for 20 years we're always finding new ways to extract data from the database and the data goes back 100 years even though, obviously, a lot of that had to be keyed in," Blacksberg said. "We have generations of information to ponder."

The records expand to thousands of burials and former owners of the property and is all localized.

Solving Problems With Revelation

With the help of Revelation and OpenInsight, the company recently solved an issue with recording employee time cards.

"I wanted to replace our time and attendance swipe card system with our own version of it," Blacksberg said. Ruane helped him design a custom touchscreen-based interface.

"We built it from scratch and it works online with our database," Blacksberg said. Now the staff has a big touchscreen that it uses to clock in for work. "And we have a touchscreen on our office computer." The solution unifies office and staff hours with the rest of the database, Blacksberg said.

Along with maintaining records, the cemetery offers a landscaping service that sends out between 3,000 and 4,000 invoices annually. This was previously done by hand but is now done by the OpenInsight software. And, noted Blacksberg, the flexibility of the solution allows the company to modify the way it does billing.

"Once those bills go out, the returning payment is keyed back into our operational management system so that work orders that have to be created are based on payment," Blacksberg said. "We'd have this bottleneck in the days before the computer system and that process took weeks. The computer system we now use takes only 2 hours to create, and bills are printed inhouse at a time convenient to our staff."

Due to Revelation's help, the company has been able to continuously develop as the world changes. As other software can become stagnant, OpenInsight continuously updates, Blacksberg noted.

While Riverside was transitioning into digital photography, Revelation was very instrumental in the process, aiding the company when needed. "We were able to add that feature of attaching a photograph to a record of the gravesite," Blacksberg said. "Documenting the state of a gravesite, both before and after work is performed on it, gives us tremendous assistance when working with customers."

While a database failure is a fear, Revelation helps Riverside avoid this threat with its Universal Driver. It has redundancy in the system and can serve the public reliably and accurately, Blacksberg explained. "Being automated the way it is, is assurance that the company can continue doing its job for the consumer base we have," Blacksberg said.

Looking Ahead

For the moment, the company is still operating in a Windows-based interface but doesn't rule out moving to the web in the future. And with OpenInsight 10 being released soon, Blacksberg is confident that Riverside will continue to use Revelation Software products for the foreseeable future.

"What makes Revelation Software powerful is that it's not only a relational database but it's a MultiValue database," Blacksberg said. "For us, it's been a real boon. Its flexibility is exceptionally powerful."

Any time the company gets an idea, Riverside can add it to the existing program without having to rework the entire system. That's a "powerful" characteristic for the company as it continues to change with the times.

*Revelation's platform is the **heart and soul** of the company that no one ever sees.*

Riverside is considering constructing a CRM platform to help it better analyze its relationship to people who are customers versus individuals who are established family members of those who are deceased.

Blacksberg said he'd also like to refine the time and attendance program with the additional ability to deliver web-based payment and streamline the process.

Though Riverside has considered moving elsewhere in the past, instead, it is mulling a decision to upgrade its software with Revelation.

Porting to another environment would be expensive and unnecessary, Blacksberg said. "One of the most difficult aspects of information automation is communicating between technology people, programmers, designers, and hardware and software people, and application end users. I would say in my experience in working with computer folks, I haven't met anyone better than the people at Revelation." ■

MV Vendors Address the Needs of an

EVOLVING MARKET IN 2018

Fresh challenges and opportunities in data management and analytics are constantly emerging. A new year brings a renewed focus on addressing challenges head on.

As we embark on 2018, some of the issues that are top-of-mind for many companies include compliance with the EU's upcoming GDPR, the need for integration with additional data management frameworks as well as cloud platforms, support for more languages, and stronger security.

In the annual MultiValue Special Report this year, *DBTA* asks: What are the challenges you see users facing and how is your MV platform addressing them? ■

—Stephanie Simone

Mike Ruane
President and CEO
Revelation Software
www.revelation.com

At Revelation Software, we perceive that our resellers, who are our users, are facing challenges in the areas of security, big data, and BYOD (Bring Your Own Device) when providing vertical applications to their end user base.

Our OpenInsight v10 product has been completely redesigned with a modern and intuitive design interface to address each of these issues.

With regard to security, we have provided the ability to encrypt the transfer of data to/from the server and client workstations. OpenInsight also provides an option to encrypt Data at Rest (DAR). The product also offers single sign-on, Open ID Connect, and a new authentication module that supports enhanced industry standard security processes.

To solve the big data issue, OpenInsight v10 introduces a Cloud-Based Filing System (CBFS) which will allow the OpenInsight toolset to operate on data stored in Couchbase, which is one of several cloud databases. A cloud database is an internet-accessible NoSQL database key/value store.

Today's vertical applications need to run on multiple platforms. OpenInsight v10 has addressed the BYOD paradigm with our OpenInsight for Web (O4W) toolset included with OpenInsight v10. O4W is a rapid application development tool that lets programmers generate responsive forms, reports, and dashboards in both desktop and mobile browser mode.

Keith Lambert
VP, Marketing and Business Development
Kore Technologies
www.koretech.com

Our clients are faced with the challenges of adopting new applications (e.g., cloud-based and mobile) into their ecosystems and integrating data across their enterprises. They must also maintain and invest in their current MultiValue ERP applications and infrastructure.

RESTful Web Services are now the standard for real-time communication between applications. However, message-based, near real-time integration is still desirable for specific projects. For example, building an enterprise data warehouse from multiple data sources and servers, or when the technology is not compatible with REST standards.

Kore Technologies is helping its clients and partners manage this hybrid integration model using our Kourier Integrator and Kourier REST solutions.

Kourier's Quick Start workbench and data profiling tools accelerate and simplify building a multi-source enterprise data warehouse. Near real-time data updates are processed using Kourier's high-performance ETL engine, powered by Kore's Net Change technology.

Kourier REST gives developers advanced features such as subscriptions (Webhooks) for creating real-time REST APIs for integrating their MultiValue application with other systems. Using a "clicks not code" approach and Kourier's API framework, developers are more productive when creating and maintaining their integrations. All requests to access the MultiValue data are secured and rated using the Kourier REST Gateway.

John Bramley
VP, Applications Platforms
Rocket Software
www.rocketsoftware.com

There have been and will be many improvements to the Rocket MultiValue (MV) Application Platform, and they all put our customers and partners first.

Our customers in Europe are facing GDPR, a broad regulation. Rocket MV provides capabilities to fulfill many GDPR requirements; however, compliance depends on applying these capabilities throughout your product design and implementation, plus procedural controls. Rocket offers a document that addresses articles of GDPR and the associated MV capabilities that help.

SPECIAL SECTION

MV SOLUTIONS

Another issue our MV customers face is hiring. That's why we're doubling down on tools that democratize MV and empower anyone to use these tools to drive business value, including Python support, opening MV databases to developers who use this popular language, helping our customers expand their hiring base.

UniData and UniVerse ship with the OpenSSL version 1.0.2h.fips library which has the latest security protocols and algorithms for SHA-2 support. Also, our Application Platforms now easily accommodate future OpenSSL updates without having to upgrade the MV Application Platform itself. Customers get one executable from Rocket when a new version of OpenSSL becomes available. OpenSSL libraries are released often and our customers need to protect their business from the latest security threats, without upgrading their Application Platform.

Doug Leupen
President and CEO
Entrinsik
www.entrinsik.com

At Entrinsik, we work with hundreds of organizations using MultiValue technology, more specifically by offering software that integrates MV data with other data sources.

New data analysis challenges are emerging every day, e.g., data governance, data volume, disparate sources, complex tools.

To stay competitive, it's becoming increasingly important for organizations to make data analytics the core of business operations and extend those capabilities outward to employees, customers, suppliers, and partners.

What's needed is a data discovery platform that provides an extensible, high-performance architecture and that simplifies data exploration and analytics without sacrificing functionality.

With the new Informer 5, organizations can quickly connect to their traditional or uncommon databases, spreadsheets, and unstructured data streams all without time-consuming warehousing or cubing.

Users can access, blend, and cleanse data in a few easy steps; then utilize data discovery to explore, analyze, and visualize information relevant to them.

Informer enables end users to easily manipulate and interact with data, collaborate with colleagues, and even create personalized homepages. Informer's crisp, modern user interface provides an easy to follow hierarchy and an immersive user experience that drastically cuts down on the learning curve associated with other data analysis products.

MV SOLUTIONS ▸▸ SPECIAL SECTION

Mark Pick
President and CEO
Pick Cloud
www.mypickcloud.com

Today, our users face the same challenges as everyone else—security, uptime, access to their system, and the need to interface with modern technologies to make sure they have access to the latest and greatest tech to keep pace with their competition.

Pick Cloud works exclusively with Google's Cloud Compute platform. It allows us to provide our customers with high-performance, scalable, virtual machines, as well as a fast, efficient, and environmentally friendly global network. Security, speed, and access are a breeze. We can fully provision a server in hours, not days, and de-provision in minutes. Users pay for what they actually use and not what they think they are going to use in the future.

We continue to offer OpenQM DBaaS (Database as a Service) which allows companies to pay as little as $15 per seat per month. This includes the OpenQM license, an AccuTerm license, as well as the infrastructure that is required. All they need is an internet connection.

The bottom line is that cloud computing enables MultiValue users to manage their mission-critical systems cost-effectively, securely, and, more importantly, worry free. They can focus on their business and contributing to their bottom line. We take care of the rest.

Patrick Payne
Chief Software Evangelist
Zumasys
www.zumasys.com

Here at Zumasys, we are finding that many customers are looking for a partner to assist them with the entire IT puzzle. New models, driven by companies such as Amazon, are revolutionizing how businesses interact with customers. Our customers are looking for help adjusting their business models to compete. We are finding that customers want to focus on their business model and have a partner to assist in the technology.

At Zumasys, we are committed to being a partner in that journey. Zumasys offers tools for other MV platforms, such as MV Connect for RESTful services, data analytics with MV Dashboard and Power BI integration, and cloud connectivity with AccuTerm.

We continue to add new features to jBASE that focus on the developer and help customers recruit new talent. With jABSE 5.7, we now offer built-in RESTful services, native object and class support, and Dynamic Files.

Zumasys also has a dedicated team of in-house developers and engineers to assist with any aspect of a customer's project. We can help our customers with everything from small, in-house projects—all the way to large cloud projects. ■

Uniware Becomes New Distributor for Zumasys

Zumasys, a provider of MultiValue software and cloud computing solutions, has appointed Uniware Pty Ltd, a national provider of MultiValue software and services, as a distributor for Australia and New Zealand.

Australia was one of the first countries to embrace MultiValue databases. Now, there are hundreds of organizations with aging databases and resource shortages which are looking for modern features and user interfaces to protect and modernize the many years of investment in their intellectual property.

"Zumasys and Uniware are totally aligned on our vision for the future of Pick," said Paul Giobbi, president and founder of Zumasys. "As an established independent software vendor (ISV) and IT/cloud provider, Uniware will help us spread modern Pick database software and services to the Australian market."

Through the agreement, Uniware will now resell Zumasys' full suite of MultiValue software products, including jBASE, OpenQM, AccuTerm, and MV Connect.

Rocket Software Receives Distinguished Award

Rocket Software, a global technology provider specializing in app modernization and systems optimization, has won the 2017 "Growth Partner of the Year" award from Ellucian, recognizing the company's technical support. Ellucian, a provider of software and services built to power higher education, selected Rocket for the award in recognition of the company's technical support for Ellucian's cloud-based solutions and for its role in helping the company transition to its current SaaS license model.

"Rocket technology, services, and support play an essential role in Ellucian's industry leading applications, whether deployed on-prem, in the cloud or SaaS," said Chris Westfall, vice president of business development at Ellucian. "And as our cloud offerings have evolved, Rocket worked with us to adjust our license model, which helped us meet customer needs in new ways. They've acted as a true partner."

Ellucian has been a long-time user of the Rocket MultiValue (MV) Application Platform, which combines an embedded database and development platform to offer the best of traditional database management with the flexibility, innovation, and performance today's applications require. ■

TRENDS ▸▸

Five Industries Turning Up the Heat on Data Science

By Peter Wang

As we get deeper into winter, temperatures may be dropping, but data science in the enterprise will remain red hot as more organizations look to harness the power of data. Data science gives organizations the insights needed to improve business outcomes across industries. With "data scientist" topping the list of best jobs in the U.S. and the number of data science positions skyrocketing, every company needs to get onboard the data train or risk getting run over by the competition. In particular, there are a few industries that are turning up the heat on data science in the coming year: aviation, cybersecurity, law enforcement, retail, and human resources. An eclectic list for sure, but all areas that are ripe for data science.

Aviation

By 2026, annual airline data generation will reach 98 billion gigabytes. This volume of data presents the opportunity to make airplanes connected devices: From fuel levels to flight routes, this could solve some of the aviation industry's top problems. They could avoid cost-draining, and potentially dangerous, concerns such as air congestion or changing flight courses based on weather patterns or overbooking. Data scientists have even gone as far to say that airplane food waste could be mitigated using passenger data. Oliver Wyman speculates that the entire flight ecosystem will be revolutionized by data science—from the airports, to the airlines, to the passengers. The biggest roadblocks that aviation is facing are monetary—many airlines are hesitant to invest upfront in the technology, but I anticipate more will come around when the savings potential becomes clear.

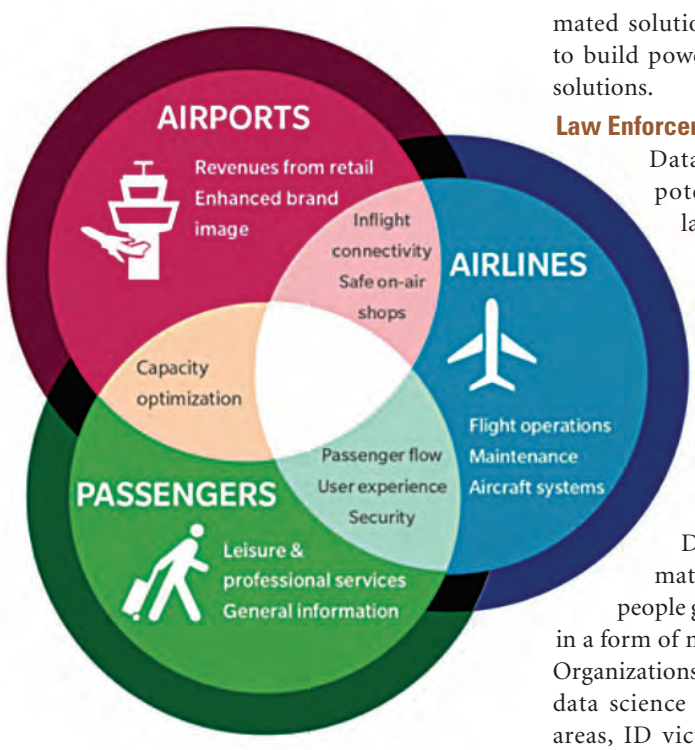

Source: Oliver Wyman Insights

Cybersecurity

I recently read an article with a shocking headline—cyberattacks could cost us more than hurricanes. Estimates show cyber threats have cost $121 billion; comparatively, Hurricane Katrina cost $108 billion in damages. While cybersecurity used to be a reactive industry, as the stakes grow higher, many organizations are looking for more proactive solutions. This is where we're seeing data science emerge, making data-driven connections that allow IT teams to spot threats before they occur. If data scientists focus on the abnormalities in a company's available data and make the move toward automated solutions, they'll be able to build powerful cybersecurity solutions.

Law Enforcement

Data science has the potential to reshape law enforcement in a number of ways. One example that's been making headlines recently is the use of machine learning to combat human trafficking. The State Department estimates 27–45.8 million people globally are trapped in a form of modern-day slavery. Organizations are starting to use data science to identify at-risk areas, ID victims, and unearth trafficking networks.

YouTube also announced recently that it's stepping up its machine learning tools to curb the distribution of terrorist videos. This comes on the heels of YouTube's partnership with Facebook, Twitter, and Microsoft that is designed to fight online terrorism. The partnership is built off of a Shared Industry Hash Database that pools data across all four companies, giving data scientists access to a wealth of information to fuel powerful machine learning tools.

Retail

Business intelligence tools have been long-standing resources for retailers, but data science is taking that to a new level. Now, more than ever before, customers

are interacting with retailers across multiple touchpoints. From smartphones to rewards programs, companies have the data needed to build solid data science tools. There are retailers using machine learning to choose the best neighborhood for their brick-and-mortar locations, to send targeted push notifications to customers, to combat online retail fraud, and to make optimized purchasing decisions.

Retail Dive has released research findings that 45% of retailers plan to start using artificial intelligence technology over the next 3 years. Its study also found that the biggest roadblocks retail marketers are struggling with are tailored recommendations, segmentation, and personalization. However, the year ahead promises a change in tune for retailers, with stores such as Amazon Go and Warby Parker leading the way.

*Businesses are changing as they **become empowered** by their own data and the **AI-backed tools** that are available to them.*

Human Resources

Over the past year, U.S. companies have spent $2 billion on HR technology, and global investments have skyrocketed. The need for data science in HR comes from a number of different places—automated interview scheduling, scanning resumes and employee reviews for keywords, improving payroll systems; the list goes on.

Additionally, there are some larger concerns for which HR is turning to data science, including the elimination of implicit bias in the hiring process. It's something that can be impossible for organizations to measure: Is every candidate receiving a fair interview? Infusing artificial intelligence into this process can alleviate that concern by automatically selecting the most qualified candidates for HR's review.

Eye on the Horizon

There is no industry, department, or job title that won't be touched by data science in the coming years. Businesses are changing as they become empowered by their own data and the AI-backed tools that are available to them. And in the months ahead, we can look forward to working with these five industries as they crank up the data science heat. ■

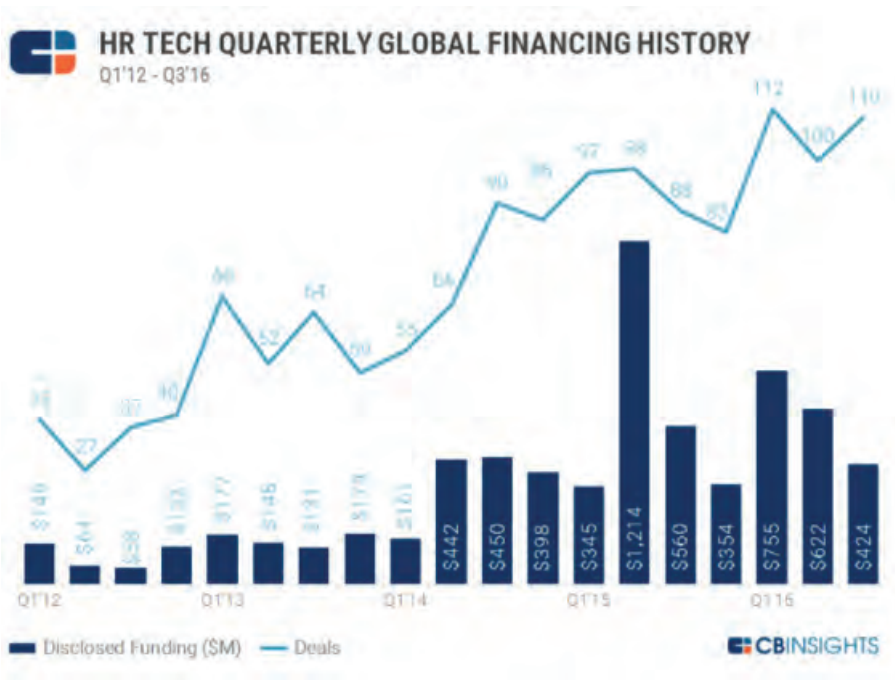

Source: CB Insights

Peter Wang is CTO and co-founder, Anaconda. Follow him on Twitter @pwang.

TRENDS ▸▸

The Enterprise Data Debt Crisis

By Eliot Knudsen

Large enterprises are facing a debt crisis. Not financial debt, but "data debt." It's a form of technical debt, and it can hamstring an organization's capacity to tackle new challenges and stifle its ability to innovate. The problem is pervasive.

A recent article in *Harvard Business Review* showed that only a shocking 3% of companies' data met basic quality standards. For years, software development teams have understood and reckoned with the future work created by making short-term tradeoffs to ship their code faster, and now IT organizations are realizing they have created massive amounts of remediation work for themselves due to decades of deprioritizing data management.

For most large enterprises, the root of this problem lies in years of treating the data generated by their operational systems as a form of exhaust rather than as a fuel to deliver great services, build better products, and create competitive advantage. Every new enterprise application deployed is essentially a data creation engine. Unless companies have a method of easily integrating each new data source to capture and leverage the new data, the debt will grow daily—and exponentially—with the number of data sources in a company. For many companies, this problem is compounded by a history of M&A, reorganizations, "data hoarding," politics, and rogue shadow IT activity.

The consequence of accumulated enterprise data debt is that many companies struggle to answer the most basic questions about their business, such as: How many customers do we have? What and how much do our customers buy? How many suppliers do we have? How much and on what do we spend with each one?

Accurate, complete, up-to-date answers to these questions require the unification of data that is often spread across dozens or even hundreds of enterprise silos. Attempts to break down the barriers between silos typically run into a host of seemingly insurmountable technical, operational, and behavioral challenges. Organizations have tried to tackle this problem for decades—often at great cost and without much success. Approaches have included application rationalization, in which they work to reduce many dozens or even hundreds of instances of an application down to a mere handful of instances; data standardization, in which IT teams try to pursue the tantalizing goal of creating a top-down comprehensive data model to meet the needs of all users, but that is hard to enforce and becomes quickly outdated; and master data management, a software-based approach that relies on developers to code up deterministic rules to integrate a few source systems and then further requires extensive manual curation of the data to create accurate master records.

These approaches have had a mixed track record historically, and with the increasing volume, velocity, and variety of data that enterprises must now manage, new options are needed. Fortunately, chief data officers (CDOs) charged with solving the data debt crisis can take a page from the playbook of their CFO colleagues.

What CDOs Can Learn From CFOs

Similar to cash, debt is a tool that managers can use to fuel business growth. But debt is also a liability—a future obligation—that eventually has to be reckoned with. If it grows out of control, the consequences are dire and sometimes life-threatening. Ask CFOs how much money their companies have and where it is, where it came from, and where it's going, and you'll get a precise, fast answer. They will have systems in place to control who has access to the company's money and what they can do with it. They will be able to show what return that money is generating, and they will be able to move that money around to get the best return for the company. In this regard, CFOs are a role model for CDOs, particularly when it comes to managing and reducing data debt.

Historically, CDOs haven't had the tools to measure, manage, and optimize their data in the way that CFOs can with cash. But that is changing. Companies such as Facebook, Google, and Apple, which treated their data as a strategic asset from their inception, have emerged to disrupt whole industries. They've built their data management infrastructure as a core capability, investing deliberately and consistently in systems to capture, store, curate, and share information. Broad (but carefully managed) access to high-quality information has fueled their rapid growth.

CDOs at large, long-standing enterprises wishing to emulate that kind of success face a more complex challenge because they must first address their accumulated data debt. But by thinking like a CFO and implementing systems and processes that manage data from creation to consumption, they can begin to create an asset that is more valuable than cash.

DataOps: A Strategy for Debt Relief

While the data debt crisis may have been decades in the making, the remedy doesn't need to take quite that long to implement. Recent advances in data management technologies, combined with approaches that borrow liberally from the success of the DevOps revolution, now offer CDOs a new

data engineering model for tackling the problem: DataOps. Just as the goal of the DevOps movement was to increase feature velocity in software, the DataOps approach seeks to radically increase analytic velocity.

As with DevOps, focus is essential to make DataOps work. CDOs should start by identifying a use case where clean, unified data is essential to achieving a high impact business outcome. Next, they should assemble a team, often called a DataOps Center of Excellence (CoE), which includes a solutions architect, line-of-business subject matter experts, and an executive sponsor.

Invariably, this group's first task is to build an inventory of their data sources relating to the identified business use case. This inventory will create a catalog of the physical data attributes and their location. This effort should start by looking at well-documented, well-understood sources (many organizations have an MDM tool, for example) and working up the data chain to find increasingly less well-governed applications and sources such as CRMs, ERPs, and data lakes. Typically, this process begins to reveal the vast scale and variety of enterprise data sources, and the business-driven use case is essential to maintain focus and avoid boiling the ocean.

The next step is to architect a blueprint for creating a data management environment. While many organizations immediately procure software at this step, they should first catalog the tools they already have before buying anything new. Most large organizations have too much software sitting undeployed on the shelf, and the most expensive mistake a fledgling DataOps CoE can make is not doing enough due-diligence at this stage. Understanding both their data inventory and software inventory will help the CoE to understand the critical capabilities they are missing. It can also illuminate why data unification projects have failed previously and enable teams to avoid repeating past mistakes.

Now, with a data inventory, a software inventory, and an understanding of key capability gaps, the CoE can architect a solution. This is when it makes sense to procure new tools in order to address the capabilities gaps. While every organization is different, there are common data curation tasks that need to be performed to create a pipeline that combines several sources to deliver high-quality, useful data. These include transformations, attribute mapping, record matching and deduplication, and classification.

It's likely that the data sources to be unified will number in the tens or hundreds, and, given the natural variety that exists across that many stores, the above operations can only scale to the challenge through automation, meaning that human intervention must be eliminated. Therefore, DataOps tools should strategically incorporate technologies such as machine learning to achieve the necessary scale and levels of automation. While machine learning for predictive analytics is unlikely to yield much value if the underlying data is of poor quality, using it to automate data preparation can yield important benefits. This is because integration models built with machine learning algorithms function as a highly transparent team of data curators and can reduce the amount of costly manual intervention by orders of magnitude.

Once harmonized, data will need to be stored, and data stores should integrate easily with downstream consumption tools such as visualization and analytics products. In a modern DataOps stack, there are typically multiple best-of-breed technologies, each performing a few tasks very well, so interoperability and the use of RESTful APIs are essential.

The ability to find data sources and unify them quickly, accurately, and at scale is the core competency of a CoE. Proficiency will allow organizations to deliver rapid, repeatable, and trusted data and analytics to end users to drive faster, better business decisions. Developing a DataOps capability can become the fastest way to pay down data debt and increase the velocity of analytics to allow organizations to build competitive advantage. While not as sexy as self-driving cars or virtual reality, in-the-trenches data engineering of this type is essential for companies that aspire to make "competing on analytics" more than a slogan.

Paying down enterprise data debt is a big challenge, and it won't happen overnight. But it's not impossible either. It's a challenge that won't be solved by any single vendor, but rather by innovative and empowered CDOs who understand that their core mission is to help their organizations manage their data as an asset and who can create the data management infrastructure to enable that change. ■

Eliot Knudsen is the data science lead at Tamr.

GERARDO DADA

In addition to being a longtime geek, **Gerardo Dada** is vice president of SolarWinds' database and applications business globally (www.solarwinds.com).

The Impact of Automation on DBAs

Are machines better equipped to do what DBAs do? That's where the first part of this series on machine learning and autonomous databases left off. Before I answer, let's first set out some of the tasks DBAs do currently (in no particular order or significance):

- Backup and recovery
- Installs, upgrades, patches
- Address downtime (due to resources, dependencies, bugs, et al.)
- Monitor health and capacity usage versus limits
- Mitigate performance issues, performance tuning
- Data modeling/design (development)
- Access control (security)/password resets
- Architecture (load balancing, high availability, and disaster recovery)
- Business intelligence
- Auditing
- Documentation
- Refreshing pre-production environments

What to Automate

Out of that list, I see a lot of tasks that would likely be better performed by a machine—a learning machine, one that makes better decisions based on more information. In fact, I have found that many issues surrounding these tasks in a prior life were caused by the human element. There's a saying, "Trust but verify," that I believe applies here. It means that you should trust the automation where it makes sense, but verify that actions being performed by automation are correct and beneficial.

Are machines better at doing some DBA tasks than DBAs? Yes. So, let automation do those tasks. I have intentionally left out which specific tasks I believe can/should be automated. The reason is that until we start to implement the automation of tasks, what may seem low-hanging fruit initially may prove quite difficult and vice versa.

What Not to Automate

Zooming out a bit, what sorts of task and decision making would be better left to humans and not automated? Here are a few where the risk/reward ratio may make them unlikely automation targets:

- Determining the value of the data
- Issues requiring out-of-the-box or more intuitive solutions
- Strategic decision making
- Tasks that could cause DBAs to lose their SQL writing skills

The last one might have gotten you to do a double-take. What I mean is that if you don't use a skill you currently have, you are likely lose it over time (perhaps not entirely, but effectively). Scripting is not similar to riding a bike, where muscle memory comes into play. Scripting is closer to math. If one doesn't use their math skills frequently, then they start to get rusty. Try factoring a complex polynomial or the derivative of a tough function. Some things about school are not missed.

Our Continually Changing World

If we can foresee a future where automation plays a more significant role and helps take on the more mundane tasks, that will free up time. How will we fill it? Here we have a few options:

- Support more instances (increase span of responsibility)
- Broaden into other IT areas (network, storage, application development)
- Go into management
- Switch career tracks
- Not have a job and let the robots do everything

I read an interesting post recently that stated 800 million jobs worldwide were at risk due to automation by 2030. The impact will not be equally distributed between job functions. In addition, different job functions will see different percentages of jobs automated. Best guess for database administrators? My guess is that around 80% of our daily/weekly tasks and activities could eventually be automated. This is purely based on personal experience and looking at tasks that would be likely to be automated.

Is that number shocking? If 80% holds true, it will certainly be impacting to the role. Should DBAs be worried? That may depend on where you are in your career and also what you believe to be the rate of change to automation (how quickly it could come to fruition). As humans, we tend to think in linear terms. However, technology tends to work in exponential terms. We may be in store for a surprise.

What impact will automation have on the DBA role? Should we start coming to terms with ML and autonomous databases? ■

GUY HARRISON

Guy Harrison, a software professional with more than 20 years of experience, is a partner at Toba Capital and the author of *Next Generation Databases* (Apress). Contact him at guy@tobacapital.com.

IOTA Aims Beyond the Blockchain

The Bitcoin bubble is a mixed bag for blockchain and cryptocurrency enthusiasts. While the incredible increase in Bitcoin's valuation has resulted in a huge windfall for early adopters and enhanced the recognition of blockchain technology, it has also highlighted the volatility of Bitcoin as a currency and the limitations of the underlying blockchain network.

Even the most hardened Bitcoin skeptic will usually acknowledge the significance of blockchain technology, while even the most ardent Bitcoin enthusiast will acknowledge the throughput and scalability limits inherent in the current blockchain implementation. During recent peaks in Bitcoin trading, transaction backlogs and transaction fees became excessive, even though the absolute transaction rates are miniscule when compared to the transaction rates supported by Visa and other payment processers.

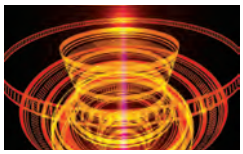

EMERGING TECHNOLOGIES

The major blockchains are all pursuing innovations that will significantly improve their ability to achieve economic and scalable transaction rates. "Off-chain" transaction processing using the Bitcoin Lightning Network or Ethereum's Raiden technology will allow small transactions to be processed independently of the main blockchain and ratified by the main chain asynchronously. Ethereum's upcoming implementation of Proof of Stake consensus mechanism should reduce the power consumption of the Ethereum network, and both Ethereum and Bitcoin blockchains are working toward sharding mechanisms which will allow the capacity of the blockchain to be increased by partitioning the network.

However, even with all of these proposed enhancements, it is completely clear that the underlying architecture of a traditional blockchain is not well-suited to the sort of massive transactional volumes that might be required to coordinate data transfers across all the internet-connected devices in the Internet of Things (IoT).

The IOTA network was envisaged as a radical departure from the traditional blockchain architecture that might meet the requirements of IoT. The architecture was designed to allow millions of IoT devices to exchange data without fees and to automatically scale in line with the size of the IoT network.

To achieve these aims, the IOTA team discarded the traditional blockchain architecture—in which each block is cryptographically dependent on all proceeding blocks—for a directed acyclic graph (DAG)—the "Tangle"—in which each transaction validates two proceeding transactions. Rather than concentrating the consensus protocol in a small number of relatively powerful "miners," IOTA distributes the Proof of Work calculations to all nodes of the network.

Furthermore, IOTA uses a quantum-resistant encryption scheme (Winternitz signatures) which is claimed to be immune to attacks by upcoming quantum computers.

The early enthusiasm around IOTA dropped sharply following some critical analysis by security experts and poor experiences by early adopters. IOTA's quantum-resistant signatures are secure only if not reused, and some users lost funds when they inadvertently reused these transaction signatures.

The IOTA team also came under criticism for overemphasising the status of partnerships with Samsung, Microsoft, and others and for inconsistent explanations for weaknesses in its hashing algorithm.

Another issue is that the existing IOTA network includes centralized coordinator nodes that are used to police transactions. The IOTA team says that these nodes can be removed when the network becomes mature, but in the meantime, the existing IOTA network is not IOTA-compliant.

The ideas underlined by IOTA are incredibly exciting. If blockchain technology is to reach its full potential, we will definitely need something like IOTA. Unfortunately, for now, it looks as if IOTA is a victim of its own popularity: Its market capitalization became huge before the network became stable, leading to disappointment and disillusionment. IOTA may or may not be the "next generation" blockchain technology we need—but the ideas and aims of the project seem very worthy. ■

CRAIG S. MULLINS

Craig S. Mullins is president of Mullins Consulting, Inc. He's an IBM Gold Consultant and the author of two best-selling books, *DB2 Developer's Guide* and *Database Administration: The Complete Guide to DBA Practices & Procedures*. Website: www.mullinsconsulting.com

It's All 'Big Data'

The term "big data" has been bandied about for a number of years now, to the point where it has been used so much that it is a part of IT culture. Hard to specifically define, yet everyone seems to have a good idea what is meant by it, big data is here to stay. And that is a good thing!

I typically define big data as being the result of a confluence of trends that coincided at the same time. Incessant data growth, alternative data storage and management systems (Hadoop, NoSQL), improved analytical tools, AI and machine learning, cloud computing, social media, sensor-based data, and mobile computing have all contributed to what we refer to as big data.

DBA CORNER

Moreover, big data refers to the shift from not just disk-based data storage, but also in-memory storage and processing. It is not just relational, but also NoSQL—and not just DBMS but also Hadoop and Spark. It is not just commercial software but also open source. Not just on-premises data and computing, but also in the cloud. Take note that these shifts are not resulting in the replacement of technology and capabilities but in the addition of it. Relational databases are not outdated or obsolete but should be a core component of your multiple data platform strategy.

Furthermore, it is not just the technology we use but how we are using it and what we are doing with it that is shifting. Big data is a result of the transition from mostly internal data to information from multiple sources; from transactional to add analytical data; from structured to add unstructured data; from persistent data to add data that is constantly on the move.

I'm sure you will recall the analyst definition of big data as consisting of four V's: Volume, Velocity, Variety, and Variability. Although interesting, and a noble attempt at defining something so all-encompassing as big data, I don't think it matters much.

Other analysts had denigrated the term big data altogether, saying that it is not about the volume of data so much as what we are doing with it. Well, sure, but that has always been the case.

To me, big data is so simple it needs no definition. It is similar to saying big dog because you immediately know what I'm talking about. Big data is all about a lot of data. Big data doesn't have to be NoSQL. And, you don't have to sit there counting up your V's to see if you're doing it. Real-time analytics on large relational data warehouses qualifies as big data to me. And, it should to you, too. Our heritage transactional systems are generating a large amount of data that is the most interesting for large enterprises to process in their big data analytics systems.

The point I'm making here is given in the title of this piece: It is all big data! And that is the way you should be thinking. How can we better store, manage, integrate, administer, analyze, and ingest all of our data to make better business decisions? How can we augment our data with partner data, social media data, and other sources of relevant data? What tools will help us do that?

If you are a DBA, then all of the management and administrative tools that you use or need to manage databases at your organization are big data tools. By adjusting the way you think about your requirements, you can focus your budget requests to hit that "big data budget" and perhaps finally get those performance or recovery tools that you've needed for years. The amount of data DBAs are managing is growing at many times the rate at which the number of DBAs is growing, so management and automation tools will be imperative to succeed.

Although I'm usually skeptical of industry trends, this one is different. Many recent IT trends have been process-orientated (e.g., object-oriented programming, web services, SOA), but I believe that data is more important than code. As I've stated before, applications are temporary, but data is forever! And if the big data trend helps us better protect, administer, and use our data, then I'm all in favor of it.

We can use the rise of big data to the forefront of computing as a means to improve data quality, institute data governance, and pay more attention to our data management infrastructure. After all, if you're going to have big data, it had better be good big data. Big data forever! ■

DAVID START
David Start is president of the Independent Oracle Users Group Board of Directors.

Which Oracle Event Should You Attend in 2018?

As the president of the Independent Oracle Users Group I often get asked for advice about which Oracle event to attend—COLLABORATE or Oracle OpenWorld. Sometimes it's just, "What is COLLABORATE?" or even, "Isn't COLLABORATE just a mini Oracle OpenWorld?" These two events are very different, and, depending on your role, one may be better for you than the other.

Oracle OpenWorld is the "big event" for Oracle. This is a very large event where Oracle showcases everything new, such as product launches, road maps, and more. You will find out where Oracle is going over the next year, what the new releases will be, and where the company is focusing. Product teams will be there to talk with customers; Oracle executives will be onsite; and some of the Oracle A list celebrities will be there. On the first day, there are usually user group sessions to get a taste of what their conferences are like and then product sessions and road maps all week long. Of course, you can't forget the entertainment at the end of the week, and usually, it's a major concert.

In my experience, the types of people that can really benefit from Oracle OpenWorld are C-level executives, management, and architects. The C-level attendees get a chance to connect with their vendors in the exhibit hall and Oracle executives in private meetings. Management attendees get a chance to see where Oracle's strategy is going so they can start planning for that. Also, if they are looking for services and service providers, the exhibit hall is a great stop. Architects get a chance to understand all of the new features as well as sit with product teams to get a voice with them.

COLLABORATE is the big Oracle users group event hosted by IOUG, Quest, and OAUG. This event is where you go to learn from users about practical implications, user case studies, and hands-on labs. COLLABORATE is focused on diving into how to use what you currently own, connecting with people who have the same issues you do, and sharing new solutions to solve them. Customers and partners of Oracle are the center of COLLABORATE. If you are looking to get up-to-speed on Oracle technologies from a user's perspective or want to dive deep into a topic with an expert, then COLLABORATE would be a good option. Just like Oracle OpenWorld, there are Oracle product teams giving road map sessions, but the big focus is customers sharing what they are doing and what works.

The types of people that can benefit most from COLLABORATE are new technologists, engineers, architects, and developers. A new technologist would have to go to many different classes over time to get the same experience he or she would at COLLABORATE and likely would not meet as many people from different industries and backgrounds.

Consider COLLABORATE to be very intense training (moving sessions about every hour). Engineers can find out how to implement many different things over the course of a week. Architects can use COLLABORATE as a way to get hundreds of new ideas to try out and to build a network of colleagues to bounce ideas off of. Developers may not want to go deep but can get a broad understanding of how to use many different Oracle technologies.

As you can see, depending on your role, both events are worth your time and money. Both serve a purpose in the community and simply allow us opportunities to be better at our jobs.

Meet me in Las Vegas to see why COLLABORATE is my favorite event of the year and to get the training, solutions, and connections you need. To find out more about COLLABORATE, visit http://collaborate.ioug.org. ■

KEVIN KLINE

Kevin Kline, a longtime Microsoft SQL Server MVP, is a founder and former president of PASS and the author of *SQL in a Nutshell*. Kline tweets at @kekline and blogs at https://blogs.sentryone.com/kevinkline.

The New SQL Server Vulnerability Assessment Tool Goes GA

The SQL Server Vulnerability Assessment tool (VA) is a feature within SQL Server Management Server (SSMS) 17.4 that scans your SQL Server instances of version 2012 and later, identifies security issues, and suggests fixes to the vulnerabilities it finds. It works for on-premises SQL Server and Azure SQL Database, whether housed on physical or virtual servers.

The tool helps ensure data privacy standards are met and adhered to, such as the EU's new GDPR legal framework; monitors database environments, where changes are frequent and difficult to track; and aids with audit or compliance reporting requirements, which can be answered with a security report.

VA runs a scan directly against a SQL Server database based upon a large knowledgebase of rules related to security vulnerabilities, security best practices, configuration recommendations, permission best practices, and awareness of sensitive data.

How Does It Work?

VA is a straightforward and easy to use tool that provides a simple way to assess a given SQL Server's database security status. You can access the feature in SSMS by selecting then right-clicking a database, then selecting Tasks ➔ Vulnerability Assessment ➔ Scan for Vulnerabilities.

SQL SERVER DRILL DOWN

- **Run the VA:** Once selected, the first thing VA requests is a path on the local file system where you would like to save the scan results. The scan is very lightweight, taking only a few seconds to run. Don't blink or you'll miss it! It is entirely read-only and will not make any changes to your database.
- **Review the Report:** Once processing is complete, VA will display a *Vulnerability Assessment Result* pane in SSMS. The report will tell you a variety of things about the security of your SQL Server database, ranging from the number of security checks performed, the number of security checks that failed, the number and type of security incidents discovered (categorized as either high-, medium-, or low-risk), and then provide detailed listings of the security checks that failed and those that were passed. The report also shows a map of sensitive data discovered within the database, including recommendations on which built-in features of SQL Server might best protect that data. (Note that when you run a VA scan on an on-premises SQL Server, you will see a richer set of security checks than when you run the scan on an Azure SQL Database. That's because there are more security checks for an on-premises SQL Server since it presents a larger surface area. A VA scan from SSMS against an Azure SQL Database will look very similar to the results you see for security in the Azure Portal).
- **Analyze the Results:** Once you've run the VA and seen the findings, it is now important to examine each failed security check to determine the extent of the security impact. The VA report provides details for an actionable remediation process for each failed check so that you are not left wondering how to resolve the issue. In addition, you can mark specific results as part of an acceptable *baseline* value. The baseline security state is basically a customization for your database security configuration, telling VA to only report on deviations from the baseline. For example, the best practice is to enable only the network protocol you are explicitly using to connect to SQL Server, usually meaning that only TCP/IP should be enabled. But let's say your application also makes use of the Shared Memory protocol. Once VA finds that contra-indication, you can accept it as part of your baseline so that it is never flagged again. Another example would be to baseline all of your CLR assemblies, since they can represent a big security backdoor. (The baseline settings are saved together with the scan result.)

You can access past VA scan files within SSMS by selecting File ➔ Open or by right-clicking the database, selecting Tasks ➔ Vulnerability Assessment ➔ Open Existing Scan.

The only shortcoming in the tool at present is that there is no easy means to automate it. However, the SSMS tools team is developing features soon for both scripting and scheduling.

Learn More

To learn more about VA and access a video tutorial against an Azure SQL Database, check out this Channel 9 demo: https://channel9.msdn.com/Shows/Data-Exposed/Track-and-remediate-potential-database-vulnerabilities-with-SQL-Vulnerability-Assessment. You can also read more documentation and directly interact with the Microsoft developers responsible for this feature at https://docs.microsoft.com/en-us/sql/relational-databases/security/sql-vulnerability-assessment. ■

TODD SCHRAML

Todd Schraml has more than 20 years of IT management, project development, business analysis, and database design experience across many industries from telecommunications to healthcare. He can be reached at TWSchraml@gmail.com.

User Solution Insufficiency Blindness

Far too often, business users seem consumed by the systems they handle. This makes them unable to define the necessary business processes or needs of the organization. All these users can do is describe what their off-the-shelf packages provide for them. In fact, most users take great pride in their "understanding" of the current system. These users often have cheat sheets that list the exact steps they must follow for one situation or another. The same users may not be able to explain why such steps are of business value; but similar to a good sorcerer's apprentice, they understand that no variances are allowed, otherwise bad things may happen.

DATABASE ELABORATIONS

Does this mean these purchased packages are perfect in every way, and meet every conceived need? No, because these acquired solutions tend to fall short. But these same users are blind to the obvious shortcomings that are evidenced by the work-arounds that become standardized within the practices and then end up leaving mangled and misused data elements all over the place in service of doing the tasks the users need to do.

Sadly, until aspects of these packages become quite unbearable, or work-arounds start failing, users won't even acknowledge there is any problem at all. The users have gone too far in accepting what they have as it is. And in so doing, how the system exists is considered the standard and therefore inevitable.

Accepting the off-the-shelf package's physical implementation as a valid description of the logical business is an easy decision for many users to make. Largely, the easiness of this acceptance consists of being able to step back, not think, and not get overly involved. It is an approach that does not rock boats, or stick necks out. And many organizations keep people burdened under enough assignments that there is no time allowed for self-reflection, so how could they possibly have time to question why?

However, these choices may leave business needs unacknowledged and unfulfilled.

And if, as with many off-the-shelf packages, a solution offers customizable components, in their desire to avoid involvement, the users' efforts will only help give rise to chaotic and contradictory customized aspects of the organization's shiny new packaged software.

Implementation of custom items relates to external results as opposed to any internal meanings and logical necessities. Reverse-engineering the results of such customizations can lead investigators to feel as if they have fallen down Alice's rabbit hole.

Under these user-blind circumstances, attempting to detect any canonical or business data models becomes problematic. These kinds of users unknowingly work against such goals. At best, these users try to explain what the developers of the solution did, as opposed to identifying exactly what is the corporate necessity. Modelers and architects must search for the few users who may see beyond the superficial workings of the existing applications. Ideally, each organization has at least one individual, a user or perhaps an executive, who rises above a given solution and can articulate the business need being solved. And, preferably these visionaries can express the future needs that are on the horizon.

A data architect must search far and wide across a great many other businesspeople to find the right individuals to help. If these useful folks are not readily available, modelers may try and work with likely candidates from the existing user pool to help them see beyond the current system. Such endeavors may seem time-consuming in the short-run, but in the long-run, they can make the difference between failure and success. These visionaries are rare gems in the corporate landscape, and they must be identified. For without the special insights from these individuals, efforts for enterprise or canonical data models are doomed to be insufficient as well. ■